



(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
27.02.2002 Bulletin 2002/09

(51) Int Cl 7: G06F 12/10

(21) Application number: 91401217.3

(22) Date of filing: 11.05.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Chauvel, Gérard
06600 Antibes (FR)

(74) Representative: Holt, Michael
Texas Instruments Ltd., EPD MS/13, 800 Pavilion
Drive
Northampton Business Park, Northampton NN4
7YL (GB)

(30) Priority: 21.08.2000 EP 00402331

(71) Applicants:

- Texas Instruments Incorporated
Dallas, Texas 75251 (US)
Designated Contracting States:
BE CH DE DK ES FI GB GR IE IT LI LU MC NL PT
SE TR AT CY
- Texas Instruments France
06271 Villeneuve Loubet Cedex, Nice (FR)
Designated Contracting States:
FR

(54) TLB with resource ID field

(57) A digital system and method of operation is provided in which several processing resources (340) and processors (350) are connected to a shared translation lookaside buffer (TLB) (300, 310(n)) of a memory management unit (MMU) and thereby access memory and devices. These resources can be instruction processors, coprocessors, DMA devices, etc. Each entry location in the TLB is filled during the normal course of action by a set of translated address entries (308, 309) along with qualifier fields (301, 302, 303) that are incorporated with each entry. Operations can be performed on the TLB that are qualified by the various qualifier fields. A command (360) is sent by an MMU manager to the control circuitry of the TLB (320) during the course of operation. Commands are sent as needed to flush (invalidate), lock or unlock selected entries within the TLB. Each entry in the TLB is accessed (362, 368) and the qualifier field specified by the operation command is evaluated (364). This can be task ID field 302, resource ID field 301, shared indicator 303, or combinations of these. Operation commands can also specify a selected virtual address entry (305). Each TLB entry is modified in response to the command (366) only if its qualifier field(s) match the qualifier(s) specified by the operation command. If a particular resource (340, 350) is not needed for a period of time, all entries in the TLB asso-

ciated with that resource are invalidated and the resource is set in a low power state in order to conserve power.


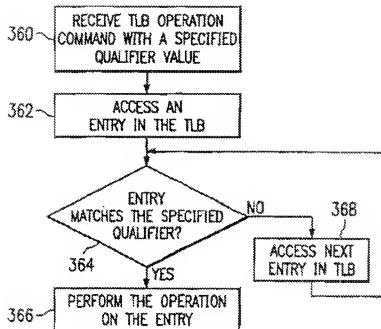


FIG. 3A

FIG. 3B



Description

[0001] This application claims priority to European Application Serial No. 00402331.3, filed August 21, 2000.

[0002] This invention generally relates to computer processors, and more specifically to improvements in translation lookaside buffers for address translation, systems, and methods of making.

[0003] Microprocessors are general purpose processors which provide high instruction throughputs in order to execute software running thereon, and can have a wide range of processing requirements depending on the particular software applications involved.

[0004] Many different types of processors are known, of which microprocessors are but one example. For example, Digital Signal Processors (DSPs) are widely used, in particular for specific applications, such as mobile processing applications. DSPs are typically configured to optimize the performance of the applications concerned and to achieve this they employ more specialized execution units and instruction sets. Particularly in applications such as mobile telecommunications, but not exclusively, it is desirable to provide ever increasing DSP performance while keeping power consumption as low as possible.

[0005] To further improve performance of a digital system, two or more processors can be interconnected. For example, a DSP may be interconnected with a general purpose processor in a digital system. The DSP performs numeric intensive signal processing algorithms while the general purpose processor manages overall control flow. The two processors communicate and transfer data for signal processing via shared memory. A direct memory access (DMA) controller is often associated with a processor in order to take over the burden of transferring blocks of data from one memory or peripheral resource to another and to thereby improve the performance of the processor.

[0006] Modular programming builds a computer program by combining independently executable units of computer code (known as modules), and by tying modules together with additional computer code. Features and functionality that may not be provided by a single module may be added to a computer program by using additional modules.

[0007] The design of a computer programming unit known as a task (or function) is often accomplished through modular programming, where a specific task is comprised of one module and the additional computer code needed to complete the task (if any additional code is needed). However, a task may be defined as broadly as a grouping of modules and additional computer codes, or, as narrowly as a single assembly-type stepwise command. A computer program may be processed (also called "run" or "executed") in a variety of manners. One manner is to process the computer code sequentially, as the computer code appears on a written page or on a computer screen, one command at a time. An alternative manner of processing computer code is called task processing. In task processing, a computer may process computer code one task at a time, or may process multiple tasks simultaneously. In any event, when processing tasks, it is generally beneficial to process tasks in some optimal order.

[0008] Unfortunately, different tasks take different amounts of time to process. In addition, the result, output, or end point of one task may be required before a second task may begin (or complete) processing. Furthermore, particularly in a multiple processor environment, several tasks may need access to a common resource that has a generally fixed capacity.

[0009] In order to better manage program tasks and physical memory, a concept of virtual memory and physical memory has evolved. Program task modules are generally compiled and referenced to virtual address. When a task is executed in physical memory, address translation is performed using a cache of translated addresses, referred to as a translation lookaside buffer (TLB). TLBs must be managed to optimize system performance as various tasks are executed.

[0010] Accordingly, there is needed a system and method for managing task processing and address translation that takes into account active tasks, active resources, and other task processing needs.

[0011] Particular and preferred aspects of the invention are set out in the accompanying independent and dependent claims. In accordance with a first embodiment of the invention, a method is provided for operating a digital system having a set of memory access resources and an associated shared translation lookaside buffer (TLB). A sequence of memory access requests is initiated by the set of memory access resources. In response to the sequence of memory access requests, a set of translated memory addresses are cached in the TLB. A resource identification value is incorporated with each translated memory address to identify which of the memory access resources requested the respective translated memory address. An operation is performed on the TLB that is qualified by the resource identification value.

[0012] In a first embodiment, an operation is performed on the TLB that invalidates only a portion of the set of translated addresses that have the selected resource identification value.

[0013] In another embodiment, the TLB has several levels, and the step of performing an operation encompasses all of the levels of the TLB.

[0014] In another embodiment, a selected resource is placed in a low power mode and its associated set of translated addresses that have the selected resource identification value are invalidated from the TLB to free unused entries.

[0015] Another embodiment of the invention is a digital system that has a translation lookaside buffer (TLB). The

TLB includes storage circuitry with a set of entry locations for holding translated values, wherein each of the set of entry locations includes a first field for a translated value and a second field for an associated resource identifier. There is a set of inputs for receiving a translation request, a set of outputs for providing a translated value selected from the set of entry locations and control circuitry connected to the storage circuitry. The control circuitry is responsive to an operation command to invalidate selected ones of the set of entry locations that have a selected resource identifier value.

[0016] In another embodiment, there is a set of resources connected to the TLB, with a set of power control circuits each connected to a respective one of the set of resources. An attribute register is connected to the set of power control circuits, and is operable to selectively control power provided to each of the plurality of resources.

[0017] Particular embodiments in accordance with the invention will now be described, by way of example only, and with reference to the accompanying drawings in which like reference signs are used to denote like parts and in which the Figures relate to the digital system of Figure 1 and in which:

Figure 1 is a block diagram of a digital system that includes an embodiment of the present invention in a megacell core having multiple processor cores;

Figure 2A and 2B together is a more detailed block diagram of the megacell core of Figure 1;

Figure 3A is a block diagram illustrating a shared translation lookaside buffer (TLB) and several associated micro-TLBs (μ TLB) included in the megacell of Figure 2;

Figure 3B is a flow chart illustrating a method of operating the TLB of Figure 3A;

Figure 4 is a block diagram of a digital system similar to Figure 1 illustrating a cloud of tasks that are scheduled for execution on the various processors of the digital system;

Figure 5 illustrates a TLB control format used to operate on the TLB and μ TLBs of Figure 3A;

Figure 6 illustrates operation of the TLB of Figure 3A for selective flushing of an entry for a given task or resource;

Figure 7 illustrates control circuitry for adaptive replacement of TLB entries in the TLB of Figure 3A;

Figure 8A is a schematic illustrating an alternative embodiment of control circuitry that utilizes a shift register for adaptive replacement of TLB entries in the TLB of Figure 3A;

Figure 8B is a schematic illustrating an alternative embodiment of the control circuitry of Figure 8A;

Figure 9 illustrates how a shared page entry is replicated for each task for different virtual address spaces;

Figure 10 illustrates how a shared page entry is used by each of the sharing tasks in a single virtual address space;

Figure 11 is a block diagram of control circuitry in the megacell of Figure 2 for dynamic control of power management systems using task attributes; and

Figure 12 is a representation of a telecommunications device incorporating an embodiment of the present invention.

[0018] Corresponding numerals and symbols in the different figures and tables refer to corresponding parts unless otherwise indicated.

Detailed Description of Embodiments of the Invention

[0019] Although the invention finds particular application to Digital Signal Processors (DSPs), implemented, for example, in an Application Specific Integrated Circuit (ASIC), it also finds application to other forms of processors. An ASIC may contain one or more megacells which each include custom designed functional circuits combined with pre-designed functional circuits provided by a design library.

[0020] Figure 1 is a block diagram of a digital system that includes an embodiment of the present invention in a megacell core 100 having multiple processor cores. In the interest of clarity, Figure 1 only shows those portions of megacell 100 that are relevant to an understanding of an embodiment of the present invention. Details of general construction for DSPs are well known, and may be found readily elsewhere. For example, U.S. Patent 5,072,418 issued to Frederick Bouteaud, et al. describes a DSP in detail. U.S. Patent 5,329,471 issued to Gary Swoboda, et al. describes in detail how to test and emulate a DSP. Details of portions of megacell 100 relevant to an embodiment of the present invention are explained in sufficient detail herein below, so as to enable one of ordinary skill in the microprocessor art to make and use the invention.

[0021] Referring again to Figure 1, megacell 100 includes a control processor (MPU) 102 with a 32-bit core 103 and a digital signal processor (DSP) 104 with a DSP core 105 that share a block of memory 113 and a cache 114, that are referred to as a level two (L2) memory subsystem 112. A traffic control block 110 receives transfer requests from a host processor connected to host interface 120b, requests from control processor 102, and transfer requests from a memory access node in DSP 104. The traffic control block interleaves these requests and presents them to the shared memory and cache. Shared peripherals 116 are also accessed via the traffic control block. A direct memory access controller 106 can transfer data between an external source such as off-chip memory 132 or on-chip memory 134 and the shared memory. Various application specific processors or hardware accelerators 108 can also be included within

the megacell as required for various applications and interact with the DSP and MPU via the traffic control block.

[0022] External to the megacell, a level three (L3) control block 130 is connected to receive memory requests from internal traffic control block 110 in response to explicit requests from the DSP or MPU, or from misses in shared cache 114. Off-chip external memory 132 and/or on-chip memory 134 is connected to system traffic controller 130; these are referred to as L3 memory subsystems. A frame buffer 136 and a display device 138 are connected to the system traffic controller to receive data for displaying graphical images. A host processor 120a interacts with the external resources through system traffic controller 130. A host interface connected to traffic controller 130 allows access by host 120a to external memories and other devices connected to traffic controller 130. Thus, a host processor can be connected at level three or at level two in various embodiments. A set of private peripherals 140 are connected to the DSP, while another set of private peripherals 142 are connected to the MPU.

[0023] Figure 2, comprised of Figure 2A and Figure 2B together, is a more detailed block diagram of the megacell core of Figure 1. DSP 104 includes a configurable cache 203 that is configured as a local memory 200 and data cache 202, and a configurable cache 204 that is configured as instruction cache 206 and a RAM-set 208, which are referred to as level one (L1) memory subsystems. The DSP is connected to the traffic controller via an L2 interface 210 that also includes a translation lookaside buffer (TLB) 212. A DMA circuit 214 is also included within the DSP. Individual micro TLBs (μ TLB) 216-218 are associated with the DMA circuit, data cache and instruction cache, respectively.

[0024] Similarly, MPU 102 includes a configurable cache 223 that is configured as a local memory 220 and data cache 222, and a configurable cache 224 that is configured as instruction cache 226 and a RAM-set 228, again referred to as L1 memory subsystems. The MPU is connected to traffic controller 110 via an L2 interface 230 that also includes a TLB 232. A DMA circuit 234 is also included within the MPU. Individual micro TLBs (μ TLB) 236-238 are associated with the DMA circuit, data cache and instruction cache, respectively.

[0025] L2 traffic controller 110 includes a TLB 240 and one or more microTLB (μ TLB) 242 that are associated with system DMA block 106, host processor interface 120b for a host connected at level two, and other application specific hardware accelerator blocks. Similarly, L3 traffic controller 130 includes a μ TLB controllably connected to TLB 240 that is associated with system host 120a at level three. This μ TLB is likewise controlled by one of the megacell 100 processors.

Memory Management Unit

[0026] At the megacell traffic controller level, all addresses are physical. They have been translated from virtual to physical at the processor sub-system level by a memory management unit (MMU) associated with each core, such as DSP core 105 and MPU core 103. At the processor level, access permission, supplied through MMU page descriptors, is also checked, while at the megacell level protection between processors is enforced by other means, which will be described in more detail later. Each MMU includes a TLB and its associated μ TLBs.

[0027] The translation lookaside buffer (TLB) caches contain entries for virtual-to-physical address translation and page descriptor information such as access permission checking, cache policy for various levels, etc. If the TLB contains a translated entry for the virtual address, the access control logic determines whether the access is permitted. If access is permitted, the MMU generates the appropriate physical address corresponding to the virtual address. If access is not permitted, the MMU sends an abort signal via signal group 244 to the master CPU 102. The master CPU is identified by the value of the R-ID field. On a slave processor such as a hardware accelerator the R-ID is equal to the R-ID of the master CPU.

[0028] Upon a TLB miss, i.e., the TLB does not contain an entry corresponding to the virtual address requested, an exception is generated that initiates a translation table walk software routine. The TLB miss software handler retrieves the translation and access permission information from a translation table in physical memory. Once retrieved, the page or section descriptor is stored into the TLB at a selected victim location. Victim location selection is done by software or with hardware support, as will be described later.

Translation Table

[0029] To provide maximum flexibility, the MMU is implemented as a software table walk, backed up by TLB caches both at the processor sub-system and megacell level. This allows easy addition of new page size support or new page descriptor information if required. A TLB miss initiates a TLB handler routine to load the missing reference into the TLB. At the Megacell 100 level, a TLB miss asserts a miss signal in signal group 244 and is routed via system interrupt router 250 to the processor having generated the missing reference or to the processor in charge of the global memory management, via interrupt signals 251, 252.

[0030] Translation tables and TLB cache contents must be kept consistent. A flush operation is provided for this reason and will be described in more detail later.

[0031] An address reference is generally located within the μ TLB or main TLB of each processor sub-system; how

ever, certain references, such as those used by system DMA 106 or host processor 120, for example, to access megacell memories can be distributed within L2 traffic controller 110 and cached into L2 system shared TLB 240. Because system performance is very sensitive to the TLB architecture and size, it is important to implement efficient TLB control commands to lock entries for critical tasks or unlock and flush those entries when a task is deleted without degrading the execution of other tasks. Therefore, each TLB and L2 cache entry holds a task-ID. Commands are supplied to flush locked or unlocked entries of a TLB/TLB corresponding to a selected task.

[0032] As part of the page descriptor information, the MMU provides cacheability and bufferability attributes for all levels of memory. The MMU also provides a "Shared" bit for each entry to indicate that a page is shared among multiple processors (or tasks). This bit, as standalone or combined with the task-ID, allows specific cache and TLB operation on data shared between processors or/and tasks. The MMU may also provide additional information, such as memory access permission and access priority as described later.

[0033] All megacell memory accesses are protected by a TLB. As they all have different requirements in term of access frequencies and memory size, a shared TLB with individual μ TLB backup approach has been chosen to reduce the system cost at the megacell level. This shared TLB is programmable by each processor. The architecture provides enough flexibility to let the platform work with either independent operating systems (OS) on each processors or a distributed OS with a unified memory management, for example

[0034] The present embodiment has a distributed operating system (OS) with several domains corresponding to each processor but only a single table manager for all processors. Slave processors do not manage the tables, in a first embodiment slave processors R-ID are equal to the R-ID of the master CPU. In another embodiment, they could, however, have a different R-ID to control their TLB entries lock/unlock entries corresponding to some of their own tasks or flush all their entries, when putting themselves in sleep mode to free entries for the others processors. Having different R-ID provides a means to increase security in a concurrent multi-processor environment, processor X can not access memory allocated to processor Y.

[0035] In another embodiment with several independent OS(s), for example, there will be independent tables. These tables can be located in a memory space only viewed by the OS that they are associated with in order to provide protection from inadvertent modification by another OS. As they manage the virtual memory and task independently, the R-ID provides the necessary inter-processor security. R-ids are managed by a single master CPU. This CPU can make TLB operations on all TLB entries. TLB operation or memory accesses from slave processor are restricted by their own R-ID. The CPU master will have rights to flush out entries belonging to another processor in a different OS domain.

[0036] The organization of the data structures supporting the memory management descriptor is flexible since each TLB miss is resolved by a software TLB-miss handler. These data structures include the virtual-to-physical address translation and all additional descriptors to manage the memory hierarchy. A list of these descriptors and their function is described in Table 2. Table 1 includes a set of memory access permission attributes, as an example. In other embodiments, a processor may have other modes that enable access to memory without permission checks.

Table 1 -

Memory Access Permission	
Supervisor	User
No access	No access
Read only	No access
Read only	Read only
Read/Write	No access
Read/Write	Read only
Read/Write	Read/Write

Table 2 -

Memory Management Descriptors	
Execute Never	provides access permission to protect data memory area from being executed. This information can be combined with the access permission described above or kept separate.
Shared	Indicates that this page may be shared by multiple tasks across multiple processor.

Table 2 - (continued)

Memory Management Descriptors	
Cacheability	Various memory entities such as individual processor's cache and write buffer, and shared cache and write buffer are managed through the MMU descriptor. The options included in the present embodiment are as follows: inner cacheable, Outer cacheable, inner Write through/write back, Outer write through/write back, and Outer write allocate. The terms Inner and outer refer to levels of caches that are built in the system. The boundary between inner and outer is defined in specific embodiment, but inner will always include L1 cache. In a system with 3 levels of caches, the inner correspond to L1 and L2 cache and the outer correspond to L3 due to existing processor systems. In the present embodiment, inner is L1 and outer is L2 cache.
Endianness	determines on a page basis the endianness of the transfer.
priority	Indicates a priority level for the associated memory address region. Memory access can be prioritized based on this priority value.

MMU/TLB Control Operation

[0037] Figure 3A is a block diagram illustrating a shared translation lookaside buffer (TLB) 300 and several associated micro-TLBs (μ TLBs) 310(0)-310(m) included in megacell 100 of Figure 2. On a μ TLB miss, the shared TLB is first searched. TLB controller 320 is alerted by asserting a μ TLB miss signal 324. In case of a hit on the shared TLB, the μ TLB that missed is loaded with the entry content of the shared TLB 300. In the case of a miss in shared TLB 300, the shared TLB alerts TLB controller 320 by asserting a TLB miss signal 326. Controller 320 then asserts an interrupt request signal 328 to system interrupt controller 250. Interrupt controller 250 asserts an interrupt to the processor whose OS supervises the resource which caused the miss. A TLB entry register 330 associated with TLB controller 320 is loaded by a software TLB handler in response to the interrupt. Once loaded, the contents of TLB entry register 330 are transferred to both shared TLB 300 and the requesting μ TLB at a selected victim location as indicated by arcs 332 and 334.

[0038] A separate TLB entry register 330 is only one possible implementation and is not necessarily required. The separate register TLB entry register is a memory mapped register that allows buffering of a complete TLB entry (more than 32 bits). A TLB value is not written directly in the TLB cache but is written to the TLB entry register first. Because of the size of an entry, several writes are required to load the TLB entry register. Loading of a TLB cache entry is then done in a single operation "Write TLB entry". Advantageously, other μ TLBs associated with other modules can continue to access the shared TLB while the TLB entry register is being loaded, until a second miss occurs. Advantageously, by controlling access to the TLB via the TLB entry register, CPUs have no direct access to TLB cache internal structure and thus the risk of partial modifications inconsistent with the MMU tables is avoided.

[0039] The sequence of operations to update a TLB cache entry after a miss is:

- 1 - the software TLB handler writes to the TLB entry register,
- 2 - the software TLB handler sends a command to write the TLB entry, which transfers a value from TLB entry register to a preselected victim TLB cache entry; and
- 3 - control circuitry checks and preselects a next victim TLB entry, in preparation for the next miss. In this embodiment, this step is generally performed in background prior to the occurrence of a miss.

[0040] Advantageously, TLB cache entries can be preemptively updated under OS software control to prevent TLB miss by pre-loading a new entry, using the following sequence of operation:

- 1 - control circuitry checks and selects a TLB entry, referred to as a victim TLB cache entry.
- 2 - the software TLB handler writes to the TLB entry register, and
- 3 - the software TLB handler sends a command to write the TLB entry, which transfers a value from TLB entry register to the selected victim TLB cache entry.

[0041] The priority on the shared TLB is managed in the same way as priority on a memory access. One or more resources can be using the shared TLB. One or more resources can program the shared TLB. The replacement algorithm for selecting the next victim location in the shared TLB is under hardware control. A victim pointer register 322 is maintained for each TLB and μ TLB to provide a victim separate pointer for each. A typical embodiment will use a round robin scheme. Different TLBs within a single megacell can use different replacement schemes. However, in an

embodiment in which the system has a master CPU with a distributed OS, this master CPU could also bypass the hardware replacement algorithm by selecting a victim entry, reading and then writing directly to the Shared TLB, for example.

[0042] In this embodiment, each shared TLB has 256 entries. Each μ TLB is generally much smaller, i.e., has fewer entries, than the shared TLB. In various embodiments, each shared TLB has 64-256 or more entries while μ TLBs generally have 4-16 entries. The penalty for a miss in a μ TLB is small since a correct entry is generally available from the shared TLB. Therefore, the present embodiment does not provide direct control of the victim pointers of the various μ TLBs; however, direct control of the victim pointer of shared TLBs, such as 212, 232, and 240, is provided.

[0043] Each entry in a TLB has a resource identifier 301 along with task-ID 302. Resource-IDs and task IDs are not extension fields of the virtual address (VA) but simply address qualifiers. Resource IDs are provided by a resource-ID register associated with each resource, such as R-ID register 342a associated with resource 340 and R-ID register 342n associated with resource 350. Resource 340 is representative of various DMA engines, coprocessor, etc within megacell 100 and/or an external host connected to megacell 100. Resource 350 is representative of various processors within megacell 100. Each resource 340, 350 typically has its own associated R-ID register; however, various embodiments may choose to provide resource ID registers for only a selected portion of the resources. A task ID is provided by a task-ID register, such as task-ID register 344a associated with resource 340 and task-ID register 344n associated with resource 350. A task register associated with a non-processor resource, such as DMA, a coprocessor, etc, is loaded with a task value to indicate the task that it is supporting.

[0044] In another embodiment, only processor resources 340, 350 that execute program modules have an associated programmable task-ID register. In this case, a system wide default value may be provided for access requests initiated by non-processor resources such as DMA. The default value may be provided by a programmable register or hardwired bus keepers, for example.

[0045] Advantageously, with the task-ID, all entries in a TLB belonging to a specific task can be identified. They can, for instance, be invalidated altogether through a single operation without affecting the other tasks. Advantageously, the resource ID permits discrimination of different tasks being executed on different resources when they have the same task number. Task-ID number on the different processors might not be related, therefore, task related operations must be, in some cases, qualified by a resource-ID.

[0046] In another embodiment, the R-ID and Task_ID registers are not necessarily part of the resource core and can be located elsewhere in the system, such as a memory mapped register for example, and associated to a resource bus. The only constraint is that a task_ID register related to a CPU must be under the associated OS control and updated during context switch. R-ID must be set during the system initialization. In some embodiments at system initialization, all R-ID and Task-ID registers distributed across the system are set to zero, which is a default value that causes the field to be ignored. In other embodiments, a different default value may be used. In other embodiments, R-ID "registers" provide hardwired values.

[0047] Referring still to Figure 3A, each TLB entry includes a virtual address field 305 and a corresponding physical address field 308 and address attributes 309. Various address attributes are described in Table 1 and Table 2. Address attributes define conditions or states that apply to an entire section or page of the address space that is represented by a given TLB entry. An S/P field 306 specifies a page size such as 64kB and 4kB for example. Naturally, the page size determines how many most significant (ms) address bits are included in a check for an entry.

[0048] Each TLB entry also includes "shared" bit 303 and a lock bit 304. All entries marked as shared can be flushed in one cycle globally. A V field 307 indicates if an associated TLB cache entry is valid. V field 307 includes several V-bits that are respectively associated with R-ID field 301 to indicate if a valid R-ID entry is present, task-ID field 302 to indicate if a valid task-ID entry is present, and virtual address field 305 to indicate if a valid address entry is present. These valid bits enable the compare logic with their associated field.

[0049] As mentioned earlier, the resource ID field and task ID field in each entry of the TLB/ μ TLB can be used to improve security. During program task execution, each transaction request is checked by the miss control circuitry of the TLB/ μ TLB to determine if the entry is allowed for a specific resource or for all resources and for a specific task or for all tasks. For example, if a request is received and a valid entry is present for the proffered virtual address but a task ID or R-ID which accompany the request does not match the corresponding valid task ID and R-ID fields of the entry, then a miss is declared. If the task ID and/or R-ID fields of the entry are marked as invalid, then they are ignored.

[0050] Figure 3B is a flow chart illustrating a method of operating the TLB of Figure 3A. As discussed above, the TLB is filled during the normal course of action by a set of translated address entries along with qualifier fields that are incorporated with each entry. As will be described in more detail below, operations can now be performed on the TLB that are qualified by the various qualifier fields.

[0051] In step 360, an operation command is received by the control circuitry of the TLB. This command is sent by the MMU manager during the course of operation. Commands are sent as needed to flush (invalidate), lock or unlock selected entries within the TLB. These operations will be described in detail later.

[0052] Step 362 accesses a first entry in the TLB and reads the qualifier field specified by the operation command.

This can be task ID field 302, resource ID field 301, shared indicator 303, or combinations of these. Operation commands can also specify a selected virtual address entry.

[0053] Step 364 compares the qualifier specified by the operation command with the qualifier field read from the TLB entry. If they match, then the operation is performed on that entry in step 366. If they do not match, then the next entry is accessed in step 368 and compare step 364 is repeated for the next entry.

[0054] Step 366 performs the operation specified in the operation command on each entry whose qualifier field(s) match the operation command. In this embodiment, the operation can invalidate an entry by resetting valid bit field 307, and lock or unlock an entry by appropriate setting of lock bit 304.

[0055] Step 368 access each next TLB entry until all entries have been accessed. In this embodiment, all μ TLBs associated with a shared TLB are also accessed as part of the same operation command.

[0056] Other embodiments may provide additional or different operations that are qualified by the qualifier fields of the present embodiment or by additional or other types of qualifier fields. For example, resource type, power consumption, processor speed, instruction set family, and the like may be incorporated in the TLB and used to qualify operations on the TLB.

[0057] Figure 4 is a block diagram of a digital system similar to that of Figure 1 illustrating cloud of tasks that are scheduled for execution on the various processors of the digital system. Typically, each software task includes a task priority value that is commonly used by an operating system to schedule an order of execution for a set of pending tasks 1440.

[0058] In this illustration, a circle such as 1442 represents a task, with a task name "c" and a task priority of 12, for example. Likewise, task 1443 has a task name "r" and a priority of 15, where a lower number indicates a higher priority. If the set of tasks 1440 are assigned to three processors, then an operating system on each processor forms a ready to execute queue, such as ready queue 1446 in which task "c" is scheduled for first execution, then task "a" and finally task "b" according to priority values of 12, 15, and 50 respectively. The Task ID register in each processor is loaded when a task is invoked.

[0059] Table 3 illustrates several portions of instruction code sequences in which a task is spawned. From line 1 to line 5, task "c" is active and spawns a new task, "audio" on line 5. The kernel is then invoked to instantiate the new task and create the associated TCB. An eight bit (numbers of bits can be more or less) task-ID field is memorised in the TCB at line 11. During the context switch (reschedule in line 13) before launching the "audio" task, the kernel loads task-ID register 1412 with the task-ID value held in the TCB (Table 4) or in another table. At line 14, the new task is now active.

Table 3 - Setting Task ID at the Start of a Task

```

1 // (Task c code execution)
2 Instruction 1
3 -----
4 instruction n
5 Taskspawn("audio",200,0,5000,(FUNCPTR)audio, // (Task c code execution:
   instruction n+2)
6 // (Kernel code execution)
7 -----
8 TaskCreate()
9 // (taskcreate code execution)
10 -----
11 SetTaskAttributeID(TID)
12 -----
13 // Kernel reschedule code execution
14 // (Task Audio code execution)
15 Instruction 1
16 -----

```

[0060] Table 4 is an example task control block that is used to define a task-ID. Typically, the OS uses a 32-bit task-ID that is in fact an address that enables the OS to locate task information (TCB). At line 4, an execution priority value is defined that is used by the operating system to schedule execution of the task. At line 5, a task-ID value is defined that is used to set the task ID register when the task is instantiated.

Table 4 - Setting Task ID Using a TCB

```

1   TCB (task control block)
2   Typedef struct TCB
3   {
4       UINT OS-priority
5       UINT Task_ID
6       ...
7   #if CPU_FAMILY == xx
8       EXC_INFO excinfo;
9       REG_SET regs;
10      ...
11  #endif
12  }

```

[0061] In other embodiments, other means than a TCB may be provided for storing the task ID.

[0062] Referring again to Figure 3A, task-ID field 302 can be set in response to information provided at line 5 of the TCB illustrated in Table 4. This information can be used directly by the MMU manager when loading a new entry in TLBs. This information could also be part of the page table descriptor in the MMU page table and loaded as part of the MMU software table walk.

[0063] In the present embodiment, task-ID information is not maintained in page tables but is inserted by the TLB miss handler at the time of a TLB fault by using the task_ID value of the transaction request that caused the TLB fault. Other embodiments may use other means for setting the task-ID field in the TLB entry, such as by storing this information in a separate table or in the MMU page tables, for example. In the present embodiment the Valid bit associated with the task-ID field is loaded through the MMU table walk and is part of the MMU tables. Thus, when the TLB miss handler accesses a page table in response to a TLB miss, it queries the task-ID valid bit field of the MMU page table; if this bit field is asserted, then the TLB miss handler asserts the task-ID valid bit in the TLB entry and loads the task-ID value from the task-ID register of the requester that caused the TLB miss into task ID field 302. If the task-ID valid bit field of the MMU page table is not asserted, then the TLB miss handler deasserts the task-ID valid bit in the TLB entry and the task-ID value from the task-ID register of the requester that caused the TLB miss is ignored.

[0064] In the present embodiment, the shared bit field 303 is loaded through the MMU table walk and is part of the MMU tables. Typically, shared pages are defined by the OS in response to semaphore commands, for example.

[0065] In another embodiment, shared bit information is not maintained in page tables but is inserted by the TLB-miss handler at the time of a TLB fault by accessing the TCB directly based on the task ID of the request that caused the fault. The TCB is located by the TLB-miss handler via a look-up table keyed to the task ID value. Other embodiments may use other means for setting the shared bit in the TLB entry by storing this information in a separate table, for example.

[0066] R-ID field 301 is set by using the R-ID of the request that caused the fault. A Master CPU could also load value in this field during the programming of a TLB entry by taking this information from the MMU tables or separate tables, for example.

[0067] Figure 5 illustrates a TLB control word format used to operate on the TLB and μ TLBs of Figure 3A in response to control operations as defined in Table 5. TLB control word format 400 includes a task-ID field 402, resource-ID field 404 and virtual address field 406. Note that the virtual address field refers to a page address, therefore lsb address bits that refer within a page are not needed. In some embodiments, certain of the processors might not be allowed to invalidate entries other than their own.

[0068] As described previously, during execution of a program, the R-ID and Task-ID field comes from a register associated with a requester during each memory system access request. In a system embodiment with multi-processors with multiple independent Operating Systems (OS), the R-ID is static and indicates which of the resources is

accessing a given location (address). The Task-ID indicates which of the tasks (or processes) of this resource is doing the access. The task ID is dynamic and changes on each context switch. For these systems, restricting operations on a system TLB to the associated resource is important to optimize the main system TLB usage. Each OS controls the TLB entries it uses.

[0069] However, another system embodiment might be controlled by middleware that supports a unified task and memory management. For those, the notion of R-ID might disappear and be treated as part of the task ID. Restriction of TLB command based on R-ID would not be necessary in those systems and the field R-ID could be re-used to extend the task-ID field. In that case, TLB control format 410 may be used in which the R_ID field is not needed. Recall that the R-ID of the requestor is provided with each transaction request, therefore control operations specified using format 410 can be confined to entries associated with the requestor.

[0070] A processor can initiate various control operations on a TLB by writing a control word conforming to appropriate format to a specific memory mapped address associated with TLB controller 320. This control word can specify a target virtual address entry and an associated task ID or an associated resource ID. Depending on the operation, unneeded fields are ignored. For example, the operation "invalidate all entries related to an R-ID" will only use the R-ID field 404. The format and type of operation can be distinguished by using different memory mapped addresses, for example. Each address corresponds to a different TLB operation. Another embodiment would be to use a different processor instruction opcode for each of the TLB operation that would drive the appropriate control signal connected to TLB controller 232. A state machine in TLB controller 320 then executes the requested control operation. These TLB control operations are listed in Table 5. These operations are described in more detail below. For many of the operations, certain processors in an embodiment will be restricted to only affecting their own entries. This restriction is enforced by using the resource-ID signals 2106 provided with each write to TLB controller 320 as part of each memory access request.

Table 5 – TLB Control Operation

Invalidate entry with VA
Invalidate all entries related to a Task-ID
Invalidate all entries related to a R-ID
Invalidate all shared entry
Invalidate all entries of a task except shared
Invalidate All entries
Lock/UnLock entry
Lock/Unlock all entries related to a task-ID/R-ID
Read TLB entry
Write TLB entry
Check and select victim TLB entry
Set victim TLB entry

[0071] In another embodiment, the control operations can be invoked by executing an instruction that invokes a hardware or software trap response. As part of this trap response, a sequence of instructions can be executed or a control word can be written to a selected address, for example. In another embodiment, one of the processors may include instruction decoding and an internal state machine(s) to perform a TLB or Cache control operation in response to executing certain instructions which may include parameters to specify the requested operation, for example.

[0072] For an "invalidate entry" operation, a Virtual page address (VA) is provided in VA field 406 of the control word and the other fields of the control word are ignored. This generates an entry invalidate operation on the corresponding virtual address entry. Note that all processors of a given megacell embodiment might not be allowed to invalidate entries others than their own. In that case, the R-ID value from the R-ID register of the requestor is used to qualify the operation.

[0073] For an "invalidate all entries related to a task" operation, all entries corresponding to the provided task identifier are invalidated. This allows a master-processor to free space from the shared TLB by invalidating all entries of a task belonging to another processor. In this case, the control word provides a task-ID value and an R_ID value. Processors other than the master-processor can free space from the shared TLB by invalidating all entries of one of its own tasks. This operation invalidates all the entries corresponding to the provided task and resource identifier or to a task of the resource requesting the operation. The R-ID value from the R-ID register of the requestor is used to qualify the operation.

[0074] For an "invalidate all entry related to a Resource" operation, all entries corresponding to RID field 404 of the control word are invalidated. Note that all processors of a given megacell embodiment might not be allowed to invalidate entries other than their own. This provides, however, the capability to a master processor to free space from the shared TLB by invalidating all entries of another processor. The R-ID value from the R-ID register of the requestor is used to qualify the operation.

[0075] For an "invalidate all shared entries" operation, all entries in the TLB marked as shared for the requester are invalidated. The R-ID register value limits the effect of this operation, as discussed above.

[0076] For an "invalidate all entries of a task except shared entries" operation, all entries in the TLB for a task specified in the control word not marked as shared for the requester are invalidated. The R-ID value from the R-ID register of the requestor limits the effect of this operation, as discussed above.

[0077] For an "invalidate all entries" operation, all entries in the TLB matching the R-ID of the requester are invalidated. For the master CPU, the operation invalidate all entry regardless of the R-ID. If all of the R-ID registers distributed in the system have the same value, then this operation invalidates all entries.

[0078] For a "lock/unlock entry" operation, a control word is written providing the VA which needs to be locked/unlocked. This operation sets or resets lock field 304 in the selected entry. Restriction on R-ID applies as above.

[0079] For a "lock/unlock all entry related to a task" operation, a control word is written providing the task identifier which needs to be locked/unlocked. Restriction on R-ID applies as above.

[0080] In the case in which an independent OS is running on each processor, each OS can initiate the above operations. In that case, these operations must be restricted to entries with a resource identifier (R-ID) belonging to the requester.

[0081] In the case of a single master OS, task and memory management can be viewed as unified, removing the need for an R-ID. The R-ID can be an extension of the task-ID. In an embodiment, in which the R-ID is hard-coded, the field R-ID in the TLB simply needs to be disabled (associated Valid bit is cleared) via a configuration control register. Disabling the R-ID is equivalent to having a single R-ID for all the system or for part of the system.

[0082] As mentioned above, a global control bit can be used in an embodiment to determine if all the above functions must be limited to the entry corresponding to the resource ID requesting the operation.

[0083] Although it is preferable to have the same page size for memory management on all processors, it is not mandatory. In a shared system, the TLB supports all page sizes of the system, in response to S/P field 306. Therefore, in a different embodiment, a TLB may support a different set of page sizes.

[0084] Table 5 also lists some additional operations that are provided which allow a software TLB handler to access the shared system TLB: Read TLB entry, Write TLB entry, Check and select victim TLB entry, and Set victim TLB entry. These are described in more detail below.

[0085] For a "Read TLB entry" operation, an entry in the TLB pointed to by the victim pointer is transferred into TLB entry register 330. The TLB entry register can then be read and analyzed by the software TLB handler. Again this operation might be restricted to the master CPU for security.

[0086] For a "write TLB entry" operation, the contents of the TLB entry register is transferred to a selected victim entry of the TLB.

[0087] The "check and select victim TLB entry" operation has multiple functions. Its first purpose is to determine an index value for the replacement of an entry. However, it can also be used to find out if an entry is already in the TLB. The R_ID & Task_ID & VA fields of a corresponding entry are checked for a match against a proffered virtual address entry. If there is no match, then the victim pointer is positioned according to the chosen replacement algorithm. This replacement can be random, cyclic, etc. The second usage is to verify if a given page is present in the TLB. If a matching entry is found, the victim entry points to this matching entry, and a flag bit in the status register is set to indicate this condition.

[0088] The "Set victim TLB entry" operation allows the software TLB handler to select a particular entry as the next victim. This is useful to support certain lock mechanisms software replacement algorithms.

[0089] As indicated earlier, each control operation is performed by a state machine within TLB control circuitry 320 in response to writing to a selected memory mapped address. For example, for the operation "invalidate all entries related to a task", all entries with a matching task-id TAG are invalidated in response to a single command, including the shared TLB and the associated μ TLBs. In the present embodiment in which the TLB is a fully associative memory, the operation can be done in one cycle or as a loop as most appropriate.

[0090] As mentioned above, control operation affect the shared TLB and the associated μ TLBs for the various operations based on task-ID, resource-ID and shared bits. In an embodiment in which both μ TLBs and TLB are fully associative, the flush and/or Lock/unlock can be done by the same command in the same cycle. But if the μ TLB is fully associative and TLB is set associative, for example, a single command is still used, but the operation into the set associative TLB will be executed entry by entry by a HW loop. This will take longer time. If both the μ TLB and TLB are fully associative there will typically be a single control block. If the μ TLB is fully associative and TLB set associative, there may be separate control blocks 320, but the same command effects all of the control blocks. Alternatively, an embodiment may require sending copies of the operation command separately to each of the separate control blocks.

[0091] Figure 6 is a simplified block diagram of the TLB of Figure 3A and will now be referred to explain selective invalidation of an entry for a given task or resource, as listed in Table 5. Processor 2100(m) is representative of one or more requestors that access TLB 2130. A physical address bus 2104(m), resource ID signals 2106(m), and task ID signals 2108(n) are provided by each processor 2100(m) for each TLB request. Traffic controller 2110 provides request priority selection and sends the highest priority request to TLB 2130 using physical address bus 2104, resource ID signals 2106, and task ID signals 2108 to completely identify each request.

[0092] A task-ID field 302 and/or a resource ID field 301 stored as independent fields in the TLB TAG array is used to selectively invalidate (flush) all entries of a given task or a given resource (requester). A state machine within control circuitry 2132 receives a directive from a processor to perform an invalidation operation, as described above. The operation directive specifies which task-ID is to be flushed using format 400 or 410 (see Figure 5).

[0093] For operations which use task ID field 402 in the control word, state machine 2132 accesses each entry in TLB 2130, examines the task-ID field, and if there is a match that entry is flushed by marking the valid bits in its valid field 307 as not valid. Thus, a single operation is provided to flush all entries of a given task located in a TLB. As discussed above, in this embodiment, the TLB cache is made of several levels of set associative TLB and μ TLB, and all levels are flushed simultaneously in response to a single operation directive command by accessing each entry sequentially in a hardware controlled loop.

[0094] For operations which use both task ID field 402 and R-ID field 404 in the control word, state machine 2132 accesses each entry in TLB 2130, examines the task-ID field and the resource ID field, and if there is a match in both the task ID and R-ID fields that entry is flushed by marking all valid bits in its valid field 307 as not valid. Advantageously, this allows discrimination of entries belonging to tasks from different resources that have the same task ID number. When the R-ID valid bit is set, an entry is not flushed if its R-ID field 301 does not match the value provided on R-ID signals 2106. This operation only invalidates entries with a valid task-ID.

[0095] In a similar manner, the selective invalidation operation "invalidate all entries related to a R-ID" is performed by examining the R-ID 301 field of each entry and if there is a match in the R-ID field that entry is flushed by marking its valid field 307 as not valid. This operation only invalidates entries with a valid R-ID.

[0096] Likewise, the selective invalidation operation "invalidate all shared entries" is performed by examining the share field 303 of each entry and if there is a match in the shared field that entry is flushed by marking its valid field 307 as not valid. All entries marked as shared can be flushed in one cycle.

[0097] In the present embodiment, when shared entries are flushed, state machine 2132 ignores the task ID field since shared page entries may be used by different tasks having different task IDs. In an alternative embodiment, shared entry flushing could also be qualified by the task ID field. Alternatively, shared entry flushing could also be qualified by the task ID field, but only if the task ID valid bit in valid field 307 is asserted indicating a valid task ID value is in field 302.

[0098] Figure 7 is a simplified block diagram of the TLB of Figure 3A and will now be referred to explain selective lock/unlocking of an entry for a given task or resource, as listed in Table 5. Advantageously, in this multi-processor system with system shared TLB, an innovative scheme of adaptive replacement is provided for controlling the TLB on a task basis, as discussed above. In order to support such a function in the most optimized way, an adaptive replacement algorithm taking into account locked entries and empty entries is provided. TLB full signal 2240 is asserted when one or more valid bits in field 307 is asserted for each TLB entry location. TLB miss signal 2242 is asserted to indicate a miss occurred in response to a transaction request from processor 2100(m), which invokes a TLB handler as described earlier.

[0099] When the TLB is full with no locked entries, pseudo-random replacement based on a simple counter (Victim CNT) 2234 is used to select the victim entry. Another embodiment would be to keep a pseudo random replacement and to check the lock bit on a miss. If it is locked, signal 2244 is asserted and the victim counter is incremented further until a non-locked entry is found. This is done automatically by the control circuitry connected to victim counter 2234 so that response time of the TLB handler routine is not impacted.

[0100] When the TLB is not full, the victim counter is incremented until an empty entry is found. This is done automatically by the control circuitry connected to victim counter 2234 so that response time of the TLB handler routine is not impacted.

[0101] After a flush entry operation is performed, the victim "counter" is updated with the location value of the flushed

entry and stays unchanged until a new line is loaded in order to avoid unnecessary searching.

[0102] An alternative implementation provides the capability to do the victim entry search instantaneously by providing in an external logic the lock and valid bit or by using a CAM, for example. In another alternative embodiment, a shift register and associated circuitry is used to point to the next location in the TLB that is either not valid or valid and not locked.

[0103] Figure 8A is a schematic illustrating an alternative embodiment of victim selection circuitry 2234 that utilizes a shift register for adaptive replacement of TLB entries in the TLB of Figure 3A. Shifter 850 is used to point to the next location in the TLB that is either not valid or valid AND not locked. Assuming the TLB has n entries, shifter 850 has only $n-1$ positions. Position zero of the TLB is reserved as a victim location when all entries are valid and all entries 1 to $n-1$ are locked.

[0104] Lock bits 804 are equivalent to lock field 304 of Figure 3A, except they are implemented as individual storage bits rather than as part of a TLB memory array so that they can be monitored by reservation circuitry comprising AND gate 854 to form a shift control signal 856 that is asserted when all of the monitored lock bits are asserted. Also, the individual lock bits can be set by individual gates 852[x] in response to a lock operation. When signal 856 is asserted, only entry location S[0] is available. If the shifter indicates S[0] and a lock request occurs, an error is signaled to the CPU by gate 852[0] because the position S[0] is reserved as unlockable. The TLB miss handler in the OS can then decide to remove one of the already locked entries in order to lock this new one.

[0105] Valid bits 807 are equivalent to the VA valid bit of valid field 307 of Figure 3A, except they are implemented as individual storage bits rather than as part of a TLB memory array so that they can be individually monitored by AND gate 858 to form an "all valid" control signal 860 that is asserted when all of the monitored valid bits are asserted.

[0106] During operation, shifter 850 has one bit set to one, such as bit S[1], and all the other bits set to zero selecting the S[1] entry as a candidate victim entry. If either the S[1] entry is not valid $V[1] = 0$ or all entries are valid and the S[1] entry is not locked (all Valid AND $V[1] = 1$ AND $L[1] = 0$), as determined by skip circuit 870[1], the shifter stops with a stop_shifter signal 874 asserted. Signal 874 is provided by OR gate 872 which receives outputs from each of a set of skip circuits 870[x] connected to each of valid bits 807 $V[1]-V[n-1]$. In this case, entry S[1] is selected as the next victim entry.

[0107] Otherwise, the shifter continues its search. S[2] is set and all the other bits of the shifter are zero. The same condition is checked for S[2] position and if true, the shifter stops on the victim entry 2. Otherwise, the shifter continues until an unlocked entry is found and selected. By doing one check per clock cycle, the shifter stops on the first sequential position it finds available for replacement.

[0108] The shifter starts (enable_clk = 1) a search after each new load entry (TLB-miss). Advantageously, the latency to find the victim entry is therefore hidden due to TLB-miss occurrence and the time required to handle the TLB-miss, which may be many CPU cycles.

[0109] Still referring to Figure 8A, an embodiment is illustrated that has one reserved unlockable entry. Other embodiments may have several ununlockable entries reserved (m). In that case, once $n-m$ entries are locked (all-locked = true), the victim entry selection iterates cyclically between 0 and $m-1$.

[0110] The position S[0] reserved in case of all_locked case is not really used, meaning that the TLB size is really $n-1$. An alternative embodiment would be to have a shifter with n location to avoid losing any entry. Lock request on unlockable entry zero (S[0]) would raise a flag and position the victim pointer on the first unlocked entry. The CPU can then read the content of the victim location and decide to use this entry to lock the desired entry. This would add potential latency on lock operation but remove the loss of TLB entries.

[0111] Another alternative implementation to avoid loss of a TLB entry is to execute the all-task-id-unlocked operation as a loop of n . In that case, a "locked-counter" can be used to detect if more than $n-m$ entries are locked and to thereby keep m entries unlocked. Every new lock entry request increments the locked-counter. The locked-counter is decremented through the all-task-id-unlocked loop.

[0112] Figure 8B is a schematic illustrating an alternative embodiment of the control circuitry of Figure 8A using reservation circuitry comprising a locked-counter 880 and comparator 882. Signal lock-auth(orized) 884 remains asserted as long as the count value is less than $n-m$. In this implementation unlock operation takes n cycles, but no TLB entry is lost. If a new lock request occurs once lock_auth is cleared, the new entry is not locked, but gate 866 asserts an error signal that can set a flag or an interrupt error can be returned to the CPU. The OS lock-error handler can then decide if another entry can be unlocked to let the new one be locked.

[0113] In this embodiment, lock bits L[n] can be part of the TLB cache memory instead of discrete logic because all lock and unlock operation are done one entry at a time (selected by the shifter). Similarly, if an up-down counter is provided to generate the all_valid signal, then the V[n] bits can also be part of the TLB memory. The skip logic can be reduced to a single set on the output of the memory and the OR 872 is removed.

[0114] Referring again to Figure 7, the function "Lock/Unlock all entries of a given task" listed in Table 5 is implemented by the comparison of the task-id field 302 of each entry in the TLB. If this field matches a task-id value 402 supplied in the control word (see Figure 5), the entry is locked by setting the associated lock bit 304 or unlocked by

clearing the associated lock bit 304 of each matching entry depending on the requested operation. In an embodiment of a TLB implemented with a memory array, the function is done through a hardware loop using a finite state machine located in control circuitry 2232, for example. In an alternative embodiment of a TLB implemented with a content addressable memory (CAM), all entries with the same task-ID can be locked or unlocked in one cycle.

[0115] As discussed above, lock/unlock request are restricted as mentioned above by the R-ID provided on signals 2105. When R-ID field 301 does not match the value provided on R-ID signals 2106 the entry is not locked/unlocked.

[0116] Thus, Lock/unlock operation on the TLB based on task-ID and optionally qualified by R-ID is provided. A pseudo-random replacement algorithm for the TLB is changed into a sequential replacement algorithm upon detecting an empty entry location or a locked victim entry location.

[0117] Figure 9 illustrates how a shared page entry is replicated for each task for different virtual address spaces. In this example, there are illustrated two tasks, referred to simply as task 1 and task 2. Each task may occupy several pages of virtual address space that are mapped to corresponding pages of physical address space. These various pages hold code, data, etc; however, for simplicity only one or two pages for each task are illustrated here for simplicity.

[0118] Task 1 has a page VA1 in virtual address region 900. Page VA1 is mapped to physical address space 910 at physical address page PA1. Task 2 has a page VA2 in virtual address region 902. Page VA2 is a shared page and is mapped to the same physical address page PA1. Also illustrated is a second page owned by task 2 in virtual address page VA3. This page is not shared and is mapped to physical address page PA2.

[0119] Table 6 illustrates six entry locations of an example TLB that is loaded with entries for task 1 and task 2. The R-ID and attribute fields are not illustrated, for simplicity. The page size field S/P holds a size value of M, but different values can be specified. Entry 2 holds page VA1 of task 1, which is shared, as indicated by the shared bit S being set to "1." Entry 4 holds page VA2 of task 2, which is shared, as indicated by the shared bit S being set to "1." Entry 5 holds page VA3 of task 2, which is not shared, as indicated by the shared bit S being set to "0." In this case, each shared entry is replicated for each task because they are in different VA regions.

[0120] Advantageously, when either task 1 or task 2 is terminated and physical memory reclaimed, all shared entries can be expunged from the TLB by performing an "invalidate all shared entries" operation directive. In this case, an invalidate all shared entries operation will invalidate entry locations 2 and 4.

Table 6 -

Example TLB for Tasks in Same VA Space									
Entry #	Task ID	Vt	S	L	VA page	V	S/P	PA page	
1						0			
2	Task 1	1	1	x	VA1	1	M	PA1	
3						0			
4	Task 2	1	1	x	VA2	1	M	PA1	
5	Task 2	1	0	x	VA3	1	M	PA2	
6						0			

[0121] Still referring to Figure 9, in an alternative system embodiment in which the operating system does not use overlaid virtual address space, each task is in its own virtual address space. In that case, region 900 and region 902 would be disjoint and separate VA spaces. In the case of overlaid virtual address space, the TLB will need to be flushed each time a context switch occurs to change the execution thread from one task to a different task. Typically, the "active user process (or task)" is mapped on a defined VA range, which is the same for all user tasks. Translation tables are modified at context switch to map the current user task in this range, and the OS tasks have another dedicated VA range whose translation does not need to be changed on context switches. In some OSes, MSBs of the VA can be used to identify the task ("process ID"), in order to reduce flushes of the TLB and caches.

[0122] In the latter case, a task ID is not needed in the TLB since each task is distinguished by the virtual address space, as illustrated Table 7. In this case, separate virtual address spaces are distinguished by mbs of the VA page, for example. In this case, each shared entry is replicated for each task because they are in different VA spaces.

Table 7 -

Example TLB for Tasks in Same VA Space						
Entry #	S	L	VA page	V	S/P	PA page
1				0		

Table 7 - (continued)

Example TLB for Tasks in Same VA Space						
Entry #	S	L	VA page	V	S/P	PA page
2	1	x	VA1	1	M	PA1
3				0		
4	1	x	VA2	1	M	PA1
5	0	x	VA3	1	M	PA2
6				0		

[0123] In the case in which the TLB must be flushed for each context switch, an operation to "flush all entries of a task except shared" is useful. For example, when a task is completed or suspended, its entries could be removed by this command, but any entries that were shared by a still active task would be spared. In another embodiment, entries for the OS could be marked shared. Advantageously, in this alternative embodiment, during a context switch, the entries relating to the OS and marked shared would not be flushed, for example.

[0124] Figure 10 illustrates how a shared page entry is used by each of the sharing tasks in a single virtual address space. In this example, there are illustrated two tasks, referred to simply as task 1 and task 2. Each task may occupy several pages of virtual address space that are mapped to corresponding pages of physical address space. These various pages hold code, data, etc; however, for simplicity only one or two pages are illustrated here for simplicity.

[0125] Task 1 has a page VA1 in virtual address region 1000. Page VA1 is mapped to physical address space 1010 at physical address page PA1. Task 2 has a page VA2 in virtual address region 1002. Page VA2 is mapped to physical address page PA2. Also illustrated is a virtual address page VA3 that is shared by task 1 and task 2. This page is mapped to physical address page PA3.

[0126] Table 8 illustrates six entry locations of an example TLB that is loaded with entries for task 1 and task 2. The R-ID and attribute fields are not illustrated, for simplicity. Entry 2 holds page VA1 of task 1, which is not shared, as indicated by the shared bit S being set to "0." Entry 4 holds page VA2 of task 2, which is not shared, as indicated by the shared bit S being set to "0." Entry 5 holds shared page VA3, as indicated by the shared bit S being set to "1." Note that the Valid Task ID (VT) bit is set to 0 to cause the task ID field to be ignored. In this case, the shared entry is used by each of the sharing tasks in the same VA space; therefore only one entry is needed and the task ID field is ignored.

[0127] Advantageously, when either task 1 or task 2 is terminated and physical memory reclaimed, all shared entries can be expunged from the TLB by performing an "invalidate all shared entries" operation directive. In this case, the invalidate all shared entries operation will invalidate entry location 5.

Table 8 -

Example TLB for Tasks in Same VA Space								
Entry #	Task ID	Vt	S	L	VA page	V	S/P	PA page
1						0		
2	Task 1	1	0	x	VA1	1	M	PA1
3						0		
4	Task 2	1	0	x	VA2	1	M	PA2
5	xxxx	0	1	x	VA3	1	M	PA3
6						0		

[0128] Figure 11 is a block diagram of control circuitry in the megacell of Figure 2 for dynamic control of power management systems using task attributes. This figure illustrates a portion of the processing system of Figure 1, showing a detailed block diagram of an autonomous processor (MPU 12), coupled to a coprocessor 16 along with other peripheral devices 100a and 100b. MPU 12 includes core circuitry 102, comprised of various core blocks 104a, 104b, and 104c. Core 102 further includes a Current Task ID register 106, a Task Priority register 108 and a Task Attributes register 110. Core 102 is coupled to a cache subsystem 112, including and instruction RAMset cache 114, a local RAM 116, an n-way instruction cache 118, an n-way data cache 120, a DMA (direct memory access) channel 122, and microTLB (translation lookaside buffer) caches 122a, 122b, and 122c. MPU 12 further includes voltage select circuitry 124 for selecting between two (or more) voltages to power the MPU 12.

[0129] The voltage select circuitry 126 provides a supply voltage to the MPU 12. As is well known in the art, the voltage needed to support processing circuitry is dependent upon several factors; temperature and frequency are two of the more significant factors. For tasks where a high frequency is not needed, the voltage can be lowered to reduce energy consumption in the processing system 10.

[0130] One or more coprocessors and other peripheral devices may be used by the MPU 12 for various functions. The coprocessor 16 is used to provide high speed mathematical computations. Peripheral A 100a could be a input/output port, for example. Peripheral B could be a pointing device interface, such as a touch screen interface.

[0131] The MPU core 102 provides the processing function for MPU 12. This processing function is broken into multiple discrete blocks 104. Each block performs a function that may or may not be needed for a given task. For example, floating point arithmetic unit, a multiplier, auxiliary accumulator, saturated arithmetic unit, count-leading-zeroes logic, and so on, could each be treated as a MPU Block 104.

[0132] The Current Task ID register 106 stores a unique identifier for the current task being executed on the MPU 12. Other autonomous processors would also have a Current Task ID register 106 and may be executing a task different from the current task executed by the MPU 12. The Task Priority register 108 associates a priority with the task. The Task Attributes register 110 stores a control word having fields which can enable/disable circuitry or configure circuitry to an optimum configuration.

[0133] The operation of the Task Attributes register 110 to enable or disable circuitry is shown in connection with Figure 11. The data stored in the Task Attributes register 110 has multiple fields which map to associated devices. For a simple on/off attribute, the field could be a single bit. Multiple bit fields can be provided for other functions, such as choosing between three or four voltages in the voltage select circuit 126.

[0134] Each of the components shown in Figure 11 as being mapped to the Task Attributes register 110 has circuitry that is responsive to a respective control field 128 in the register. For the voltage select circuit 126, one of multiple voltages is selected based on the value of the respective field 128. In Figure 11, Vdd0 could be chosen if the field is a "0" and Vdd1 could be chosen if the field is a "1". For a voltage select circuit with four possible voltages, Vdd0 could be chosen if the field is a "00" and Vdd1 could be chosen if the field is a "01", Vdd2 could be chosen if the field is a "10" and Vdd3 could be chosen if the field is a "11".

[0135] Each of these devices has an associated power switching circuit that supplies power to the component responsive to the value of the associated field in Task Attributes register 110. For example, coprocessor 16 can be as disabled (power off), along with peripheral A 100a, while peripheral B 100b is enabled. Disabling power to a component that is not used in a task can significantly reduce the overall power consumed by the processing system 10. Similarly, MPU block A 104a and MPU block C 104c are enabled, while MPU block B 104b is disabled.

[0136] In some cases, a hardware resource may be coupled to multiple autonomous processors. For example, a Level 2 shared memory may be coupled to both the MPU and the DSP. In cases where a hardware resource is shared between two or more autonomous processors, the resource can be coupled to the Task Attributes register 110 of each processor, and the subsystem can be enabled or disabled based on a logical operation on the associated bit values. For example, assuming that a bit value of "1" represented an "on" state for the hardware subsystem, a logical OR operation on the task attribute bits would enable the resource if either processor was executing a task that needed the resource.

[0137] As described earlier, each TLB and μ TLB entry includes a field identifying a processing resource or memory access requestor (R_ID) that "owns" that entry. This resource ID field is part of the TLB TAG array to enable requestor-selective operations, such as flushes. Each resource that can request memory access via a TLB or μ TLB has a resource ID register, such as R-ID register 1130 associated with MPU core 102, R-ID register 1132 associated with co-processor 16 and R-ID register 1134 associated with peripheral 100a. In this embodiment, peripheral 100b is a slave device, so it does not have a resource ID register.

[0138] Using the task attribute register as shown in Figure 11 can significantly reduce the power consumed by the processing system 10 by disabling circuitry which is not used by a specific task. Advantageously, when a resource is disabled because it is not needed for a specific task, any and all TLB entries associated with that resource can be flushed from the multilevel TLB by performing an "invalidate all entries related to a R-ID" for the resource ID of the disabled resource, as described earlier. This frees up entries in the TLB and thereby improves processing performance.

Digital System Embodiment

[0139] Figure 12 illustrates an exemplary implementation of an example of such an integrated circuit in a mobile telecommunications device, such as a mobile personal digital assistant (PDA) 10 with display 14 and integrated input sensors 12a, 12b located in the periphery of display 14. As shown in Figure 12, digital system 10 includes a megacell 100 according to Figure 1 that is connected to the input sensors 12a,b via an adapter (not shown), as an MPU private peripheral 142. A stylus or finger can be used to input information to the PDA via input sensors 12a,b. Display 14 is connected to megacell 100 via local frame buffer similar to frame buffer 136. Display 14 provides graphical and video

output in overlapping windows, such as MPEG video window 14a, shared text document window 14b and three-dimensional game window 14c, for example.

[0140] Radio frequency (RF) circuitry (not shown) is connected to an aerial 18 and is driven by megacell 100 as a DSP private peripheral 140 and provides a wireless network link. Connector 20 is connected to a cable adaptor-modem (not shown) and thence to megacell 100 as a DSP private peripheral 140 provides a wired network link for use during stationary usage in an office environment, for example. A short distance wireless link 23 is also "connected" to ear piece 22 and is driven by a low power transmitter (not shown) connected to megacell 100 as a DSP private peripheral 140. Microphone 24 is similarly connected to megacell 100 such that two-way audio information can be exchanged with other users on the wireless or wired network using microphone 24 and wireless ear piece 22.

[0141] Megacell 100 provides all encoding and decoding for audio and video/graphical information being sent and received via the wireless network link and/or the wire-based network link.

[0142] It is contemplated, of course, that many other types of communications systems and computer systems may also benefit from the present invention, particularly those relying on battery power. Examples of such other computer systems include portable computers, smart phones, web phones, and the like. As power dissipation and processing performance is also of concern in desktop and line-powered computer systems and micro-controller applications, particularly from a reliability standpoint, it is also contemplated that the present invention may also provide benefits to such line-powered systems.

[0143] Fabrication of the digital systems disclosed herein involves multiple steps of implanting various amounts of impurities into a semiconductor substrate and diffusing the impurities to selected depths within the substrate to form transistor devices. Masks are formed to control the placement of the impurities. Multiple layers of conductive material and insulative material are deposited and etched to interconnect the various devices. These steps are performed in a clean room environment.

[0144] A significant portion of the cost of producing the data processing device involves testing. While in wafer form, individual devices are biased to an operational state and probe tested for basic operational functionality. The wafer is then separated into individual dice which may be sold as bare die or packaged. After packaging, finished parts are biased into an operational state and tested for operational functionality.

[0145] The digital systems disclosed herein contain hardware extensions for advanced debugging features. These assist in the development of an application system. Since these capabilities are part of the megacell itself, they are available utilizing only a JTAG interface with extended operating mode extensions. They provide simple, inexpensive, and speed independent access to the core for sophisticated debugging and economical system development, without requiring the costly cabling and access to processor pins required by traditional emulator systems or intruding on system resources.

[0146] As used herein, the terms "applied," "connected," and "connection" mean electrically connected, including where additional elements may be in the electrical connection path. "Associated" means a controlling relationship, such as a memory resource that is controlled by an associated port. The terms assert, assertion, de-assert, de-assertion, negate and negation are used to avoid confusion when dealing with a mixture of active high and active low signals. Assert and assertion are used to indicate that a signal is rendered active, or logically true. Deassert, de-assertion, negate, and negation are used to indicate that a signal is rendered inactive, or logically false.

[0147] While the invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various other embodiments of the invention will be apparent to persons skilled in the art upon reference to this description. For example, in another embodiment, the TLB may be limited to a single processor and not shared, or it may include only a single level without μ TLBs.

[0148] In another embodiment, the TLB may be controlled by other means than a state machine controller, such as directly by an associated processor, for example.

[0149] In another embodiment, there may be several distinct MMUs with associated TLBs, wherein certain of the TLBs may include aspects of the invention and certain others may not.

[0150] It is therefore contemplated that the appended claims will cover any such modifications of the embodiments as fall within the true scope and spirit of the invention.

Claims

1. A method of operating a digital system having a plurality of memory access resources and an associated shared translation lookaside buffer (TLB), comprising the steps of:

initiating a plurality of memory access requests from each of the plurality of memory access resources;
caching a plurality of translated memory addresses in the TLB responsive to the plurality of memory access requests;

incorporating a resource identification value with each translated memory address to identify which of the plurality of memory access resources requested the respective translated memory address; and performing an operation on the TLB that is qualified by the resource identification value.

2. The method according to Claim 1, wherein the step of performing an operation comprises invalidating only a portion of the plurality of translated addresses that have the selected resource identification value.

3. The method of Claim 2, further comprising the step of placing a selected resource in a low power mode after invalidating all of the plurality of translated addresses that have the selected resource identification value.

4. The method according to Claim 1, wherein each memory access request includes a virtual address and a resource identification value and wherein the step of performing an operation comprises the steps of:

selecting a translated memory address cached in the TLB in response to a memory access request; and comparing the resource identification value included with the memory access request to a resource identification value incorporated with the selected translated memory address and indicating a miss in the TLB if they are not the same.

5. The method according to any preceding Claim, wherein the TLB has several levels, and wherein the step of performing an operation encompasses all of the several levels of the TLB.

6. The method according to any previous Claim, further comprising the step of incorporating a second qualifier value with each translated memory address; and wherein the step of performing an operation on the TLB is qualified by both the resource identification value and the second qualifier value.

7. A digital system having a translation lookaside buffer (TLB), the TLB comprising:

storage circuitry with a plurality of entry locations for holding translated values, wherein each of the plurality of entry locations includes a first field for a translated value and a second field for an associated resource identifier value,

a set of inputs for receiving a translation request;

a set of outputs for providing a translated value selected from the plurality of entry locations; and

control circuitry connected to the storage circuitry, wherein the control circuitry is responsive to an operation command to invalidate selected ones of the plurality of entry locations that have a selected resource identifier value.

8. The digital system of Claim 7, wherein the digital system further comprises a second level TLB connected to the TLB, the second level TLB comprising:

second level storage circuitry with a plurality of entry locations for holding translated values, wherein each of the plurality of entry locations includes a first field for a translated value and a second field for an associated resource identifier value; and

wherein the control circuitry is connected to the second level storage circuitry, the control circuitry being responsive to an operation command to invalidate selected ones of the plurality of entry locations in the second storage circuitry which have the selected resource identifier value, such that qualified entry locations in the TLB and in the second level TLB are invalidated in response to a single operation command.

9. The digital system according to any of Claims 7-8, further comprising:

a plurality of resources connected to the TLB;

a plurality of power control circuits, each connected to a respective one of the plurality of resources; and

an attribute register connected to the plurality of power control circuits, operable to selectively control power provided to each of the plurality of resources.

10. The digital system of any of Claims 7-9 being a personal digital assistant, further comprising:

a processor (CPU) connected to the TLB and thereby connected to access a memory circuit;
a display, connected to the CPU via a display adapter;
radio frequency (RF) circuitry connected to the CPU; and
an aerial connected to the RF circuitry.

5

10

15

20

25

30

35

40

45

50

55




FIG. 2A

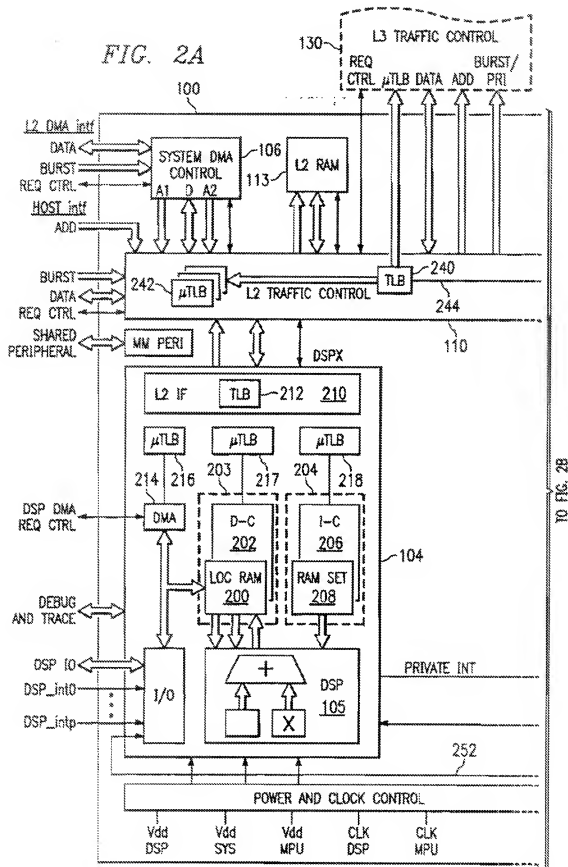
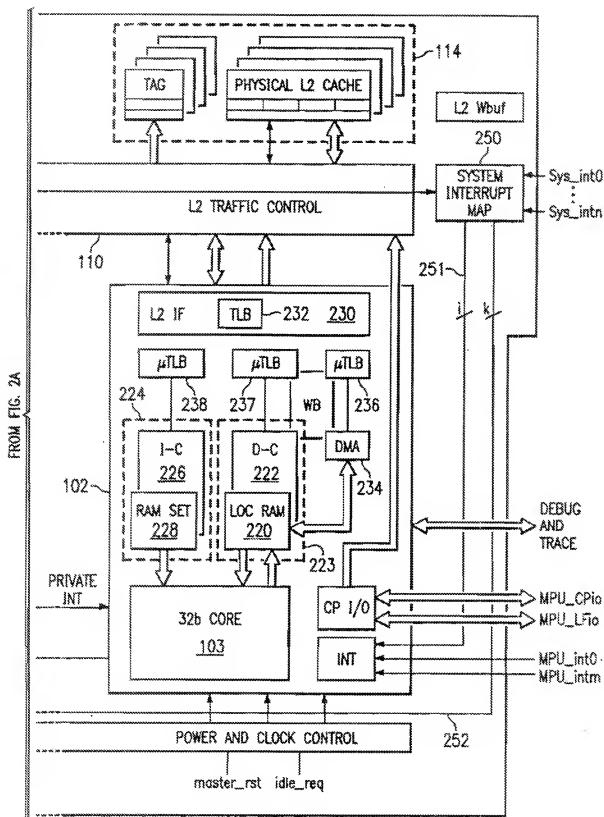


FIG. 2B






FIG. 3A

FIG. 3B

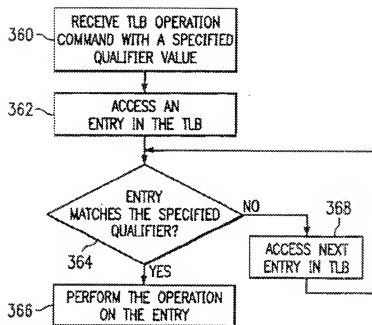




FIG. 5






FIG. 4




FIG. 6





FIG. 7






FIG. 8B





FIG. 9

FIG. 10






FIG. 11



(12) EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
04.08.2004 Bulletin 2004/32

(51) Int Cl.7: G06F 12/10

(43) Date of publication A2:
27.02.2002 Bulletin 2002/09

(21) Application number: 01401217.3

(22) Date of filing: 11.05.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Chauvel, Gérard
06500 Antibes (FR)

(30) Priority: 21.08.2000 EP 00402331

(74) Representative: Hoff, Michael
Texas Instruments Ltd.,
EPD MS/13,
800 Pavilion Drive
Northampton Business Park,
Northampton NN4 7YL (GB)

(71) Applicants:

- Texas Instruments Incorporated
Dallas, Texas 75251 (US)
Designated Contracting States:
AT BE CH CY DE DK ES FI GB GR IE IT LI LU MC
NL PT SE TR
- Texas Instruments France
06271 Villeneuve Loubet Cedex, Nice (FR)
Designated Contracting States:
FR

(54) TLB with resource ID field

(57) A digital system and method of operation is provided in which several processing resources (340) and processors (350) are connected to a shared translation lookaside buffer (TLB) (300, 310(n)) of a memory management unit (MMU) and thereby access memory and devices. These resources can be instruction processors, coprocessors, DMA devices, etc. Each entry location in the TLB is filled during the normal course of action by a set of translated address entries (308, 309) along with qualifier fields (301, 302, 303) that are incorporated with each entry. Operations can be performed on the TLB that are qualified by the various qualifier fields. A command (360) is sent by an MMU manager to the control circuitry of the TLB (320) during the course of operation. Commands are sent as needed to flush (invalidate), lock or unlock selected entries within the TLB. Each entry in the TLB is accessed (362, 368) and the qualifier field specified by the operation command is evaluated (364). This can be task ID field 302, resource ID field 301, shared indicator 303, or combinations of these. Operation commands can also specify a selected virtual address entry (305). Each TLB entry is modified in response to the command (366) only if its qualifier

field(s) match the qualifier(s) specified by the operation command. If a particular resource (340, 350) is not needed for a period of time, all entries in the TLB associated with that resource are invalidated and the resource is set in a low power state in order to conserve power.


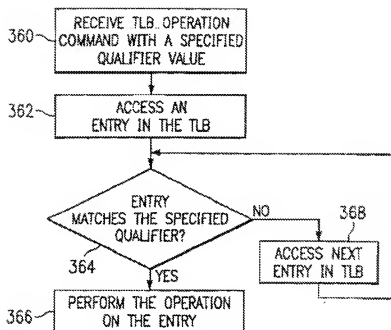


FIG. 3A

FIG. 3B





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 40 1217

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	US 5 437 016 A (IKEGAYA HIROSHI ET AL) 25 July 1995 (1995-07-25)	1, 2, 4, 6, 7	G06F12/10
Y	* figure 1 * * figure 2 * * figure 8 * * figure 9 * * column 1, line 14 * * column 3, line 3 - line 58 * * column 4, line 25 - line 38 * * column 6, line 56 - line 58 * * column 8, line 55 - column 9, line 42 *	3, 9	
Y	US 5 623 647 A (MALTRA AMIT K) 22 April 1997 (1997-04-22) * column 2, line 39 - column 3, line 19 *	3, 9	
			TECHNICAL FIELD SEARCHED (Int.Cl.7)
			G06F
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 10 June 2004	Examiner Jardon, S
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date O : document cited in the application L : document cited for other reasons X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background C : non-written disclosure P : intermediate document	

EP 01 40 1217 A3 (PCT/JP00/000000)

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 01 40 1217

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as mentioned in the European Patent Office EOP file on the European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

10-06-2004

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5437016	A	25-07-1995	JP 5012126 A	22-01-1993
US 5523547	A	22-04-1997	NONE	

LUND 15/06/2004

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82